

Article

Regime-Dependent Predictability of Cryptocurrency Distress: Cross-Sectional Evidence from Two Exchanges

Huda Aldhahi *  and Abdulrahman Alsamaani 

Department of Finance, College of Business, Imam Mohammad Ibn Saud Islamic University (IMSIU),
Riyadh 11564, Saudi Arabia; aalsamaani@imamu.edu.sa

* Correspondence: hisaldhahi@imamu.edu.sa

Abstract

Can the distress of a cryptocurrency be predicted from its market behavior, and is that predictability reliable when it matters most? Using daily data for 609 USD-quoted coins traded on Kraken between 2013 and 2025, we built a survivorship-inclusive coin-quarter panel and model the onset of severe, sustained price distress—a deep, non-recovering draw-down relative to a trailing peak. A panel logit confirmed that realized volatility, illiquidity, weak momentum, and asset youth predict distress, with a coin-stratified cross-validated out-of-sample AUC of about 0.68. Our central contribution was to show that this predictability is regime-dependent. Interactions between coin-level signals and contemporaneous market-wide volatility are jointly significant (likelihood-ratio $p < 0.001$), and a rolling-origin evaluation reveals prospective accuracy swinging from no better than chance (AUC 0.43) to strong (0.79) across years. This regime-dependence is robust across alternative distress thresholds, regime proxies, data frequencies, cluster-bootstrap inference, and replication on a second exchange (Binance), though the individual signal channels are not. Testing the most natural mechanism—rising cross-asset co-movement in turbulent markets—we find no support. Microstructure-based early-warning signals for crypto distress are thus conditionally reliable: informative in calm markets but unreliable in the turbulent conditions where warning is most valuable.

Keywords: cryptocurrency; financial distress; market microstructure; early warning; regime-dependence; panel logit

JEL Classification: G17; G32; C58; G01



Academic Editors: William Forbes,
Ann-Ngoc Nguyen and Xuan Minh
Thuy Nguyen

Received: 23 June 2026

Revised: 14 July 2026

Accepted: 1 August 2026

Published: 7 August 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Cryptocurrency markets are defined as much by disappearance as by innovation. A large share of digital assets that begin trading with active markets eventually lapse into illiquidity, are delisted, or cease trading altogether, often stranding investors in positions they can no longer exit. Whether this deterioration is foreseeable from observable market behavior—or arrives abruptly and idiosyncratically—is a question of direct consequence for investors, exchanges, and regulators in a market that remains largely outside the protections of conventional finance. This paper asks not only whether distress can be predicted, but whether that predictability holds in the conditions where it would be most useful.

A growing literature establishes that asset-level failure in crypto markets is, at least partially, predictable from market data. Recent studies apply survival models, discrete-choice specifications, and machine-learning classifiers to forecast coin death, exchange

delisting, or transition to a non-traded state, grounding their predictors in liquidity and market-microstructure theory. This work consistently finds that liquidity deterioration, elevated volatility, weak performance, and asset youth precede failure, echoing the classical corporate-distress literature in which market-based variables—returns, volatility, and size—predict bankruptcy and delisting (Shumway, 2001; Chava & Jarrow, 2004; Campbell et al., 2008). The age effect in particular is well established: Lansky (2020), studying over 2500 coins, found that younger cryptocurrencies are markedly more likely to be delisted.

We depart from this literature in two ways. First, in framing: rather than predicting outright death or delisting—a rare event on a well-curated exchange, and already the focus of several recent studies—we model the onset of severe, sustained price distress, a more frequent and economically meaningful state of deterioration that does not require a coin to vanish from the market. Second, and more importantly, is the central question. The existing predictive literature estimates a single model over a fixed sample and reports its accuracy, implicitly treating the relationship between market signals and distress as time-invariant. We ask instead whether that relationship is stable or conditional on the prevailing market regime—a question that, to our knowledge, has not been examined for cryptocurrency distress, despite ample evidence that volatility and predictability in these markets are themselves strongly time-varying. Section 2.5 positions this question explicitly against the most closely related studies.

Our answer is that predictability is decidedly conditional. Although we confirm that standard microstructure indicators predict distress onset on average—reproducing the established finding—we show that this predictability depends sharply on market-wide volatility. Coin-level signals lose their discriminating power when the entire market is turbulent, and a year-by-year out-of-sample evaluation reveals prospective accuracy ranging from no better than chance to strong. The model, in short, is most reliable precisely when it is least needed. This regime-dependence is our central contribution, and it is robust: it survives alternative definitions of distress, alternative regime proxies, a change of data frequency, and conservative cluster-bootstrap inference.

The paper makes three contributions. First, using a broad cross-section of 609 coins and the full available history of a single major exchange—and, by including coins that ultimately become inactive, avoiding the survivorship truncation that Ammann et al. (2022) show distorts inference in crypto samples—we confirm out-of-sample that microstructure indicators predict distress. Second, we establish that this predictability is regime-dependent, with jointly significant interactions between coin-level signals and market-wide volatility, and we quantify the resulting instability through rolling-origin evaluation. Third, we delineate the boundaries of the finding: the aggregate regime-dependence is robust across specifications, but the decomposition into individual channels is not, and the most natural mechanism—rising co-movement in turbulent markets—is tested and rejected. Honest demarcation of what is and is not robust is, we argue, part of the contribution.

The remainder of the paper is organized as follows. Section 2 reviews the related literature and positions our contribution. Section 3 describes the data and Section 4 the methodology. Section 5 presents the baseline, out-of-sample, regime, and mechanism results; Section 6 reports an extensive battery of robustness checks. Section 7 discusses implications and limitations, and Section 8 concludes.

2. Literature Review

Our study draws on, and contributes to, four literatures: the prediction of cryptocurrency failure and distress from market data, the documentation of crypto attrition and survivorship bias, the study of time-varying volatility and connectedness, and the emerging theme of regime-dependent or conditional predictability. We review each, identify the

studies closest to our own, and state precisely where our contribution lies and where it does not.

2.1. Predicting Asset Failure and Distress from Market Data

The premise that asset failure is foreshadowed by observable market variables originates in the corporate-distress literature, where hazard models built on returns, volatility, and size predict bankruptcy and delisting (Shumway, 2001; Chava & Jarrow, 2004; Campbell et al., 2008). This logic has been transported to cryptocurrency markets by a fast-growing body of work that applies survival models, discrete-choice specifications, and machine-learning classifiers to coin death, exchange delisting, or transition to a non-traded state. An early formulation is due to Fantazzini and Zimin (2020), who recast the credit risk of a crypto-coin as a “probability of death”—a state in which its price collapses and it becomes illiquid—and propose a zero-price-probability estimator for it; Fantazzini (2022) extends this to more than two thousand coins observed between 2015 and 2020, comparing credit-scoring, machine-learning, and time-series models across forecast horizons and finding that the best model depends on the chosen death definition and on coin age, whereas lagged trading volumes and online attention remain consistently informative. Closest to our baseline is the work on predicting “zombie” assets. Będowska-Sójka et al. (2026) ground their predictor choice explicitly in liquidity and market-microstructure theory and forecast which coins will become temporarily untradable; using predictors derived from returns, volume, and market capitalization for assets listed at least 210 days between 2015 and 2022, they report that tree-based models—random forests in particular—attain roughly 84% out-of-time balanced accuracy at a 28-day horizon, with persistent declines in trading volume among the strongest signals of impending untradability. Recent deep-learning approaches reach similar conclusions: Konuk and Güvenir (2025), classifying coin “death” from daily prices and volumes using a stratified five-fold cross-validation framework akin to ours, report areas under the ROC curve in the 0.68–0.78 range—comparable to our own out-of-sample accuracy. A parallel line targets exchange-level delisting directly: Yang and Kwon (2026), building explainable machine-learning models for Binance delistings, combined quantitative price–volume indicators with qualitative signals drawn from real-time news and social-media activity. At the level of entire trading venues rather than individual coins, Fantazzini and Calabrese (2021) modelled the probability that a crypto exchange closes, identifying cybersecurity grades, developer-team transparency, exchange age, and the number of listed coins as the main predictors. The consistent message of this strand—that deteriorating liquidity, elevated volatility, weak performance, and youth precede failure—is one our baseline results reproduce. Confirming it is the starting point of our analysis, not its contribution.

2.2. Survival, Attrition, and Survivorship Bias

A parallel strand documents the scale of crypto attrition and its consequences for inference. Lansky (2020), analyzing over 2500 cryptocurrencies in a survival-analysis framework, established a strong age effect: a newly listed coin faces roughly a 35% probability of delisting within its first year, but the conditional probability of survival rises sharply with age, so that coins trading for several years face markedly lower hazards. Gandal et al. (2021), tracking more than a thousand coins and several hundred tokens, found that about 44% of publicly traded coins are abandoned at least temporarily and documented that many are subsequently revived, complicating any binary notion of “death”—an abandonment-and-resurrection dynamic the same research team documented earlier, reporting that 71% of abandoned coins are eventually revived (Feder et al., 2018). Ammann et al. (2022), using 3904 cryptocurrencies over 2014–2021, quantified the survivorship and delisting bias that

distorts performance measurement when samples are restricted to survivors. This strand bears directly on our design: by building the panel from the full available history of all USD-quoted coins on the exchange—including those that ultimately become inactive—we mitigated the truncation these studies warn against, and our cross-exchange replication deliberately contrasts a survivorship-inclusive Kraken sample with a survivorship-biased Binance one. This strand also anchors the interpretation of our age coefficient, whose protective sign reproduces Lansky’s central finding rather than constituting a new result.

2.3. Volatility Regimes and Connectedness

A third literature establishes that cryptocurrency volatility, and its transmission across coins, is strongly time-varying and intensifies during market stress; the volatility process itself is well captured by GARCH-type models (Katsiampa, 2017). Building on the generalized variance-decomposition connectedness framework of Diebold and Yilmaz (2012) and its frequency-domain extension by Baruník and Křehlík (2018), Antonakakis et al. (2019) apply a time-varying-parameter factor-augmented VAR to the major coins and find that total connectedness exhibits large dynamic variability and that the degree of connectedness depends on the prevailing uncertainty in cryptocurrency markets, strengthening in turbulent periods. A complementary body of regime-switching work shows that crypto return and volatility dynamics differ across latent market states: Maciel (2021) found that Markov-switching GARCH specifications capture distinct low- and high-volatility regimes in the leading coins, and Figà-Talamanca et al. (2021) documented regime switches and time-varying commonalities across the cryptocurrency asset class. This literature motivates our regime variable: if market-wide conditions reshape the cross-asset volatility structure, they may also reshape the mapping from coin-level signals to distress. We draw on it for motivation rather than as a competing approach, since it characterizes the volatility process itself rather than the predictability of distress.

2.4. Conditional and Regime-Dependent Predictability

Closest to our central contribution is a nascent literature recognizing that predictability in crypto markets is conditional on market state. That return predictability is not constant but varies with the market state is well established in asset pricing: aggregate equity-return predictability is concentrated in recessions and is largely absent in expansions (Henkel et al., 2011; Rapach et al., 2010), and predictive relationships are captured more accurately once their coefficients are allowed to vary over time (Dangl & Halling, 2012; Pettenuzzo et al., 2014). Two recent crypto studies are particularly relevant. First, Buse et al. (2025), predicting value-at-risk across 105 of the largest cryptocurrencies with generalized random forests, find that the relevance of individual predictors is itself regime-dependent—long-horizon measures of standard deviation that are informative in stable periods cease to be relevant predictors in unstable, volatile ones—an out-of-sample finding that closely parallels ours from a risk-forecasting angle. Second, Lee and Ryu (2026) document state-dependence in the relationship between crypto returns and credit spreads, with the effect strengthening in high-volatility periods and persisting after controlling for broad market returns. These studies establish that state- and regime-dependence are real features of crypto markets. They differ from our work in object: the first concerns the relevance of VaR predictors, the second the return–credit-spread relationship. More broadly, Bianchi et al. (2023) document that the predictability of cryptocurrency returns is itself time-varying, with both the relevant predictors and the forecasting model shifting over the sample. To our knowledge, the regime-dependence of out-of-sample predictability of cryptocurrency distress onset—whether the very signals that flag vulnerable coins remain reliable across market conditions—has not been examined. That specific question is what we address.

2.5. Positioning and Contribution

Against this background our contribution is specific and bounded. We do not claim to be first to predict crypto distress from microstructure variables (the zombie-asset and probability-of-death literature, from [Fantazzini and Zimin \(2020\)](#) to [Będowska-Sójka et al. \(2026\)](#), does so), nor first to document state-dependence in crypto ([Buse et al. \(2025\)](#) and [Lee and Ryu \(2026\)](#) do so in adjacent settings), nor that the age or liquidity effects are novel (they reproduce [Lansky \(2020\)](#) and the microstructure literature). Our contribution is to bring these strands together on a previously unexamined question: we show that the out-of-sample predictability of distress onset is itself regime-dependent—reliable in calm markets, unreliable in turbulent ones—that this conditioning is statistically robust across distress definitions, regime proxies, data frequencies, inference methods, and a second exchange, and that the most natural mechanism for it can be tested and rejected. The honest demarcation of what is robust (the aggregate effect) from what is not (the individual channels) is part of that contribution.

Table 1 summarizes the studies closest to ours and makes the distinction explicit: the failure- and distress-prediction literature estimates a single, time-invariant model and reports its accuracy; the regime-switching literature models the volatility process itself rather than the predictability of distress; and the conditional-predictability studies concern adjacent objects. The regime-dependence of the out-of-sample predictability of distress onset—the question of this paper—is examined by none of them.

Table 1. Positioning relative to the most closely related studies.

Study	Object Predicted	Approach	Regime-Dependence of Predictability Examined?
Fantazzini and Zimin (2020) ; Fantazzini (2022)	Coin “death”	ZPP estimator; credit-scoring, ML, time-series	No
Będowska-Sójka et al. (2026)	“Zombie” (untradable) state	Tree-based machine learning	No
Konuk and Güvenir (2025)	Coin death	Deep learning	No
Yang and Kwon (2026)	Exchange delisting (Binance)	Explainable ML; news and social signals	No
Buse et al. (2025)	Value-at-Risk	Generalized random forests	Yes—relevance of VaR predictors
Lee and Ryu (2026)	Return–credit-spread relation	State-dependent regression	Yes—different object
This paper	Onset of severe, sustained distress	Panel logit with regime interactions; rolling-origin out-of-sample evaluation	Yes—the central question

3. Data

3.1. Primary Data: Kraken

Our primary data are drawn from the Kraken exchange’s published historical OHLCVT (open, high, low, close, volume, trades) records ([Kraken, n.d.](#)), which provide daily candlestick data for every traded pair from the inception of each market through 31 December 2025. Kraken is among the longest-established regulated exchanges, and its records extend back to 2013, giving an unusually long cross-section. We use the daily files for all pairs quoted in fiat US dollars. After excluding stablecoins and fiat-denominated tokens (e.g., USDT, USDC, EURT), whose distress dynamics differ fundamentally from those of speculative assets, the primary sample comprises 609 cryptocurrencies. Crucially, the sample is survivorship-inclusive: it retains coins that ultimately become inactive on the exchange rather than conditioning on survival, mitigating the truncation bias documented by [Ammann et al. \(2022\)](#).

3.2. Panel Construction

From the daily series we construct a coin-quarter panel. For each coin-quarter we compute the quarterly return, realized volatility (the annualized standard deviation of daily log returns), average daily dollar volume, average daily trade count, a quarter-over-quarter volume trend, and coin age. Returns are computed only on genuinely traded days—those with positive volume and at least one trade—to avoid spurious zero returns from inactive intervals. To ensure measurement reliability, we retain only coin-quarters with at least 30 trading days and coins with at least 90 days of history, and we winsorize heavy-tailed predictors at the 1st and 99th percentiles. The resulting estimation sample contains 2579 at-risk coin-quarters and 529 distress-onset events (20.5%), distributed across 2017–2025 (Figure 1).




Figure 1. Distress-onset events over time (quarters with adequate cross-sections, 2021 onward).

Table 2 reports descriptive statistics, including the mean of each predictor split by whether distress follows in the next quarter. The univariate patterns are economically sensible and anticipate the multivariate results: coin-quarters preceding distress exhibit higher realized volatility (1.22 versus 1.03), younger age (2.0 versus 2.6 years), and lower trading activity (1600 versus 2080 average daily trades) than those that remain healthy. The Amihud (2002) illiquidity measure is highly right-skewed—its mean (49.6) far exceeds its median (0.83)—which, together with its near-collinearity with dollar volume, motivates our use of dollar volume as the primary liquidity proxy.

Table 2. Descriptive statistics.

Variable	Mean	SD	Median	Distress	Healthy
Realized volatility	1.07	0.41	1.00	1.22	1.03
Log dollar volume	11.8	2.41	11.6	11.8	11.8
Quarterly return	0.281	0.854	0.042	0.213	0.299
Volume trend	0.223	1.07	0.114	0.279	0.209
Age (years)	2.49	1.99	2.10	2.01	2.62
Avg. daily trades	1980	4990	461	1600	2080
Amihud illiquidity	49.6	831	0.83	21.0	57.0

Note: The final two columns report the mean of each variable for coin-quarters that are versus not followed by distress onset. N = 2579 coin-quarters; 529 distress events.

Bars show the number of distress-onset events; the line shows the event rate. Onset clusters in distinct episodes rather than tracking a single market crash, and the largest waves do not coincide with the highest market-wide volatility—part of why the regime-dependence operates at the cross-sectional level rather than as a simple crash-year effect.

3.3. Replication Data: Binance

For the cross-exchange replication (Section 6.7) we use the Binance exchange’s public daily klines, covering 79 USDT-quoted coins from 2017 through 2026. Binance is the largest cryptocurrency exchange by volume, and its markets are quoted in the USDT stablecoin rather than fiat US dollars; the replication therefore varies both the venue and the quote asset. Because the Binance sample is constructed from currently-listed pairs, it is—unlike the Kraken sample—survivorship-biased toward survivors. This contrast is deliberate: a finding that holds in both a survivorship-inclusive and a survivorship-biased sample, on two different exchanges and two different quote regimes, is more credibly a property of the market than of any single dataset.

3.4. Scope and Limitations of the Data

Distress is observed on exchange, not globally: a coin that becomes inactive on Kraken (or Binance) may continue trading elsewhere, so our event captures distress as manifested on a major venue rather than global extinction. We regard this as an economically meaningful signal—delisting or dormancy on a major exchange is itself consequential for investors—and the cross-exchange replication directly addresses the concern that any single venue’s dynamics are idiosyncratic. The two samples also differ in quote asset (USD versus USDT) and survivorship profile, differences we treat as features of a demanding replication rather than as confounds to be eliminated.

4. Methodology

4.1. Notation and Panel Structure

Let $i = 1, \dots, N$ index cryptocurrencies and q index calendar quarters. The unit of observation is the coin-quarter (i, q). For each coin-quarter we observe daily open, high, low, close, volume, and trade-count data, from which we construct quarterly predictors and a binary outcome. Daily log returns are defined as follows:

$$r_{i,t} = \ln(P_{i,t}/P_{i,t-1}) \quad (1)$$

where $P_{i,t}$ is the closing price of coin i on day t . Returns are computed only on genuinely traded days (positive volume and at least one trade) to avoid spurious zero returns from inactive intervals. The full set of variables is summarized in Table 3.

Table 3. Variable definitions.

Variable	Definition	Role
RVol(i, q)	Realized volatility: annualized SD of daily log returns in the quarter	Predictor
LDVol(i, q)	Log of mean daily dollar volume (price $\times$ volume) in the quarter	Predictor
Ret(i, q)	Compounded quarterly log return	Predictor
VTrend(i, q)	Change in log dollar volume vs. previous quarter	Predictor
Age(i, q)	Days since the coin’s first observed trade, at quarter-end	Predictor
MVol(q)	Market regime: cross-sectional mean of RVol across coins in q	Moderator
Distress(i, q)	1 if coin i is in a deep, non-recovering drawdown in q	State
Onset($i, q+1$)	1 if coin i transitions into distress in $q + 1$ (the outcome)	Outcome

4.2. Predictor Construction

For each coin-quarter we compute five predictors. Realized volatility is the annualized standard deviation of daily log returns,

$$RVol_{i,q} = \sqrt{(D_q^{-1} \sum_{t \in q} (r_{i,t} - r_{i,q}^-)^2)} \times \sqrt{365} \quad (2)$$

where D_q is the number of trading days in quarter q and $r_{i,q}^-$ the mean daily return. Liquidity is captured by the log of average daily dollar volume,

$$LDVol_{i,q} = \ln(D_q^{-1} \sum_{t \in q} P_{i,t} V_{i,t}) \quad (3)$$

where $V_{i,t}$ is the daily traded volume. Momentum is the compounded quarterly return, $Ret_{i,q} = \exp(\sum_{t \in q} r_{i,t}) - 1$; the volume trend is the first difference of log dollar volume, $VTrend_{i,q} = LDVol_{i,q} - LDVol_{i,q-1}$; and $Age_{i,q}$ is the number of days between the coin's first observed trade and the quarter-end. Heavy-tailed predictors are winsorized at the 1st and 99th percentiles to limit the influence of extreme observations.

4.3. Defining Distress and Its Onset

We measured deterioration relative to a trailing peak rather than an all-time peak. Let the trailing peak be the maximum close over the preceding 365 days,

$$M_{i,t} = \max_{\tau \in [t-364, t]} P_{i,\tau} \quad (4)$$

and the drawdown $DD_{i,t} = P_{i,t}/M_{i,t} - 1$. A coin is classified as distressed in quarter q if its drawdown breaches a threshold θ both at its worst point in the quarter and at quarter-end—capturing a deep decline that does not recover within the quarter:

$$Distress_{i,q} = 1\{\min_{t \in q} DD_{i,t} \leq -\theta \text{ and } DD_{i,q(end)} \leq -\theta\} \quad (5)$$

Measuring against a trailing peak is essential: an all-time-peak definition mechanically and permanently classifies most coins as distressed, since few reclaim prior highs (in our data this inflates the unconditional distress rate above 80%). The modelled event is the onset of distress—a transition from a non-distressed to a distressed state:

$$Onset_{i,q} = 1\{Distress_{i,q} = 1 \text{ and } Distress_{i,q-1} = 0\} \quad (6)$$

The estimation sample is the at-risk set: coin-quarters that are currently non-distressed ($Distress_{i,q} = 0$), have at least 30 trading days in the quarter, and belong to coins with at least 90 days of history. The outcome is next-quarter onset, $Onset_{i,q+1}$, ensuring all predictors are known as of quarter-end and the design is free of look-ahead bias. The primary threshold is $\theta = 0.70$; $\theta \in \{0.60, 0.80\}$ are examined in robustness.

The primary threshold has an economic rationale. With annualized realized volatility averaging above 100% in our sample (Table 2), drawdowns of 30–50% lie within ordinary fluctuation for cryptocurrencies and cannot distinguish impairment from routine volatility; a useful threshold must be deep enough to separate distress from noise yet not so deep as to capture only outright collapse. At $\theta = 0.70$, sustained through quarter-end, an investor has lost more than two-thirds of value relative to the trailing-year peak with no within-quarter recovery—a deterioration of the order documented for coins that are subsequently abandoned or delisted (Feder et al., 2018; Gandal et al., 2021), and consistent with the price-collapse characterization of coin death in Fantazzini and Zimin (2020). Empirically, $\theta = 0.70$ yields an unconditional onset rate of 20.5%, positioned between the 31.0% obtained

at $\theta = 0.60$, where the event begins to absorb ordinary bear-market declines, and the 12.6% at $\theta = 0.80$, where it is dominated by terminal collapse see Section 6.1.).

4.4. Baseline Empirical Model

We estimated a pooled panel logit for the probability of next-quarter distress onset. Let $x_{i,q}$ denotes the vector of standardized coin-level predictors $\{RVol, LDVol, Ret, VTrend, Age\}$ and $\gamma_{y(q)}$ a full set of year fixed effects absorbing market-wide conditions. The model is as follows:

$$\Pr(Onset_{i,q+1} = 1 | x_{i,q}) = \Lambda(\alpha + \beta' x_{i,q} + \gamma_{y(q)}) \quad (7)$$

where $\Lambda(z) = (1 + e^{-z})^{-1}$ is the logistic CDF. Coefficients are reported as odds ratios, $\exp(\beta)$. Because coin-quarters are clustered within coins, standard errors are computed using a cluster-robust covariance estimator with clustering on coin i :

$$\hat{V}_{cl}(\hat{\beta}) = (X'WX)^{-1} (\sum_i u_i u_i') (X'WX)^{-1} \quad (8)$$

where u_i is the sum of the score contributions for coin i . The model fit is summarized by McFadden's pseudo- $R^2 = 1 - \ell_1 / \ell_0$, the ratio of fitted to null log-likelihoods.

Three considerations motivate the pooled specification over fixed-, random-, or dynamic-panel alternatives. First, it is the discrete-time hazard model of Shumway (2001): a pooled logit with time-varying covariates, time effects, and within-unit clustering—the standard framework in the distress-prediction literature on which we build. Second, coin fixed effects are unattractive in this setting: conditional fixed-effects estimation discards every coin whose outcome does not vary, a large share of the panel, whereas unconditional fixed effects suffer the incidental-parameters bias in short panels; more fundamentally, our object is out-of-sample prediction for coins never seen in estimation (Section 4.6), for which coin-specific intercepts are unavailable by construction at prediction time. Third, unobserved coin heterogeneity is partially absorbed by age and the liquidity variables, and all inference is robust to within-coin dependence through coin-clustered and, in Section 6.4, cluster-bootstrap standard errors. As a direct check, we re-estimated the regime-interaction model of Section 5.4 under an exchangeable within-coin correlation structure (GEE) and as a coin random-intercept logit: the interaction estimates are essentially unchanged—each coefficient moves by less than 0.02 across the three estimators—indicating that unobserved coin-level heterogeneity does not drive the regime-dependence result. The design is, finally, dynamic in the relevant sense: the at-risk set conditions on the lagged distress state, so the estimated hazard is state-dependent by construction.

A related question is endogeneity. Volatility and trading volume are equilibrium outcomes, jointly determined with incipient deterioration, and omitted fundamentals—project failure, security incidents, regulatory action—plausibly drive both the predictors and the outcome. Three features of the design bound, though they cannot eliminate, this concern. All predictors are measured in quarter q , whereas the outcome is onset in quarter $q + 1$, so the model is strictly predictive and free of contemporaneous feedback; the at-risk set excludes coins already in distress, so predictors are not mechanically contaminated by the distress state itself; and our claims are forecasting claims rather than causal ones—we ask whether market signals predict distress and whether that predictability is stable, not whether volatility causes distress. The coefficients reported below should accordingly be read as predictive associations, and the regime-dependence result, which concerns the stability of the predictive mapping rather than its structural interpretation, is unaffected by this caveat. We return to this limitation in Section 7.5.

All analyses were performed in R, version 4.4.1 (R Foundation for Statistical Computing, Vienna, Austria); the machine-learning benchmarks in Section 6.5 and the supplemen-

tary evaluation metrics were computed in Python 3.12 (statsmodels 0.14, scikit-learn 1.8, xgboost 3.3).

4.5. Regime-Dependence

To test whether predictability depends on market conditions, we defined a market-wide volatility regime for each quarter as the cross-sectional mean of coin realized volatilities,

$$MVol_q = N_q^{-1} \sum_i RVol_{i,q} \quad (9)$$

standardized across quarters. We then augmented the baseline with interactions between the key predictors and the contemporaneous regime:

$$\Pr(\text{Onset}_{i,q+1}) = \Lambda(\alpha + \beta' x_{i,q} + \delta MVol_q + x_{i,q}' \varphi MVol_q + \gamma_{y(q)}) \quad (10)$$

The vector φ captures how each coin-level signal's marginal effect varies with the market regime; a non-zero φ implies regime-dependent predictability. We assessed the joint significance of the interactions with a likelihood-ratio test of the augmented model (10) against the baseline (7),

$$LR = 2(\ell_{\text{int}} - \ell_{\text{base}}) \sim \chi^2(k) \quad (11)$$

where k is the number of interaction terms, complemented by the Akaike Information Criterion for model comparison.

4.6. Out-of-Sample Evaluation

Because in-sample fit overstates predictive value, we evaluated performance in two ways, both summarized by the area under the ROC curve (AUC), equivalent to the probability that a randomly chosen distress case receives a higher predicted score than a randomly chosen non-case. First, coin-stratified five-fold cross-validation: coins are partitioned into five folds, the model is trained on four folds and scores the held-out fold, and a single AUC is computed over the pooled out-of-fold predictions. Partitioning on coins (rather than observations) ensures no coin appears in both training and test data, preventing the model from memorizing coin-specific behavior. Second, a rolling-origin scheme respects the time ordering: for each year Y , the model is trained on all coin-quarters through year Y and evaluated on year $Y + 1$, yielding a sequence of genuinely prospective AUCs that reveal the stability of predictive accuracy across market conditions. Year fixed effects are omitted from the out-of-sample models, as future-year dummies cannot be extrapolated.

4.7. Mechanism Test

To probe why predictability varies with the regime, we tested the natural hypothesis that coins co-move more strongly in turbulent markets, compressing the cross-sectional dispersion that coin-level signals exploit. For each quarter we computed the average pairwise correlation of daily returns across actively-traded coins,

$$\rho^-_q = (n_q(n_q - 1)/2)^{-1} \sum_{i < j} \rho_{ij,q} \quad (12)$$

where $\rho_{ij,q}$ is the correlation between coins i and j over quarter q and n_q the number of qualifying coins (≥ 40 daily observations, non-zero variance; quarters with fewer than 10 coins are excluded). We then relate ρ^-_q to the volatility regime $MVol_q$ by Pearson and Spearman correlation. A positive, significant association would support the co-movement channel; its absence rejects it.

5. Results

5.1. Sample Composition

The final at-risk sample comprises approximately 2580 coin-quarters drawn from 609 cryptocurrencies, of which 529 are followed by a distress-onset event in the subsequent quarter—an unconditional event rate of roughly 20%. Onset events are distributed across the full 2017–2025 window rather than concentrated in a single episode, though they are more numerous in the high-issuance years of 2024–2025 (when the cross-section is largest), a feature the year fixed effects absorb. Descriptively, coin-quarters preceding distress exhibit higher prior realized volatility, lower dollar volume and trade counts, weaker quarterly returns, and younger age than those that remain healthy; univariate tests confirm that volatility ($p < 10^{-17}$), age ($p < 10^{-13}$), and dollar volume ($p < 0.001$) differ most sharply between the two groups.

5.2. Baseline Predictability

Table 4 reports the baseline panel logit (Equation (7)) with year fixed effects and coin-clustered standard errors. The estimates confirm that distress onset is predictable from market-microstructure variables, in line with prior failure-prediction studies, and identify realized volatility as the dominant signal. A one-standard-deviation increase in prior realized volatility raises the odds of next-quarter distress by a factor of 2.48 ($p < 0.001$), by far the largest marginal effect in the model. Momentum is protective: stronger quarterly returns are associated with markedly lower distress odds (odds ratio 0.75, $p = 0.001$). Age is likewise protective—older coins are more resilient—with an estimated per-year odds ratio of 0.83 ($p < 0.001$), consistent with the established finding that younger coins fail more readily. The volume trend enters positively (odds ratio 1.17, $p = 0.029$), indicating that coins experiencing recent volume expansion are, conditional on the other covariates, somewhat more prone to subsequent distress. The model achieves a McFadden pseudo- R^2 of 0.18, respectable for a cross-sectional discrete-outcome model of this kind.

Table 4. Baseline distress-onset panel logit (year fixed effects, coin-clustered standard errors).

Predictor	Coef.	Odds Ratio	<i>p</i> -Value
Realized volatility	0.908	2.48	<0.001
Log dollar volume	0.054	1.06	0.044
Quarterly return	−0.283	0.75	0.001
Volume trend	0.154	1.17	0.029
Age (per year)	−0.192	0.83	<0.001
Observations (coin-quarters)	2579		
Distress-onset events	529		
McFadden pseudo- R^2	0.18		
Akaike information criterion	2182		
Year fixed effects	Yes		

A note on the liquidity channel. The sign of the dollar-volume coefficient warrants explicit comment. In univariate comparisons and in reduced specifications that exclude correlated liquidity measures, higher dollar volume is associated with lower distress odds, consistent with a protective role for liquidity. In the full specification reported in Table 4 its coefficient is positive but only marginally significant (odds ratio 1.06, $p = 0.044$). This sensitivity reflects the strong overlap among liquidity-related variables—dollar volume, trade count, and the Amihud illiquidity ratio are mechanically related and highly collinear (a difficulty in crypto liquidity measurement documented by Brauneis et al. (2021), who find the Amihud ratio poorly correlated with benchmark measures precisely because of its mechanical link to trading activity)—so that their individual coefficients are not separately

well-identified. We therefore interpret the liquidity evidence at the level of the construct (liquidity is broadly protective) rather than reading the precise sign of any single proxy in the saturated model.

Table 5 substantiates this interpretation with explicit diagnostics. The upper rows of the table report the correlation structure of the predictors together with the alternative liquidity proxies: log dollar volume, log trade count, and log Amihud illiquidity are near-collinear, with pairwise correlations of 0.95, -0.95 , and -0.91 . Variance inflation factors tell the same story. Among the five predictors of Table 4 all VIFs are below 1.8—the estimated model itself is well-conditioned—but adding the trade count and the Amihud ratio inflates the dollar-volume VIF from 1.2 to 21.4, far beyond the conventional threshold of 10. The final rows trace the consequence: the dollar-volume coefficient swings from negative in a univariate specification to positive in the saturated model and to strongly negative when all three proxies enter jointly. The sign of any single liquidity proxy in a saturated specification is therefore uninformative, and the construct-level reading given above is the appropriate one. The reduced specifications in the final rows of Table 5 make the protective direction of liquidity directly visible: the dollar-volume coefficient is negative both in the univariate specification and, strongly so, when the full set of liquidity proxies enters jointly.

Table 5. Collinearity diagnostics for the liquidity block.

	RVol	LDVol	Ret	VTrend	Age	logTrades	logAmihud	VIF-1	VIF-2
RVol	1.00	-0.02	0.43	0.42	-0.32	0.03	0.24	1.42	2.32
LDVol	-0.02	1.00	0.19	0.20	0.36	0.95	-0.95	1.24	21.4
Ret	0.43	0.19	1.00	0.64	-0.10	0.22	-0.06	1.80	1.81
VTrend	0.42	0.20	0.64	1.00	-0.09	0.24	-0.06	1.79	1.86
Age	-0.32	0.36	-0.10	-0.09	1.00	0.35	-0.41	1.30	1.31
logTrades	0.03	0.95	0.22	0.24	0.35	1.00	-0.91	—	11.6
logAmihud	0.24	-0.95	-0.06	-0.06	-0.41	-0.91	1.00	—	20.3
Log dollar volume across specifications.									
Specification	Coefficient (per SD)							<i>p</i> -Value	
Univariate (year FE only)	-0.091							0.166	
Full, excluding volume trend	$+0.128$							0.042	
Full specification (Table 4)	$+0.101$							0.110	
Full + trade count + Amihud illiquidity	-0.671							0.011	

Note: Pearson correlations on the estimation sample and variance inflation factors (VIF-1: five-predictor model of Table 4; VIF-2: all liquidity proxies included). The final four rows report the coefficient on log dollar volume (per SD) across specifications; all specifications include year fixed effects with coin-clustered standard errors.

5.3. Out-of-Sample Performance

In-sample fit overstates predictive value, so the model's practical content rests on the out-of-sample evidence. Coin-stratified five-fold cross-validation—in which no coin appears in both training and test folds—yields a pooled AUC of 0.678. This indicates a modest but genuine ability to rank coins by distress risk on assets never seen in estimation, comfortably above the 0.50 chance benchmark, though well short of the near-perfect accuracy that would signal leakage. Because the ROC AUC can flatter a classifier under class imbalance, we complement it with imbalance-robust measures on the same out-of-fold predictions: the precision–recall AUC is 0.31, roughly 1.5 times the no-skill baseline given by the 20% event rate; the Brier score (0.154) improves on an uninformative base-rate forecast (0.161); and, classifying at a threshold equal to the training event rate, precision is 0.31, recall 0.63, and F1 0.42. A decile calibration plot (Figure 2) shows predicted probabilities tracking observed onset rates closely through the bulk of the distribution, with modest over-prediction only in the highest-risk decile.




Figure 2. Calibration of out-of-fold predicted probabilities (coin-stratified five-fold cross-validation), by decile of predicted probability. The dashed line indicates perfect calibration. Markers show deciles of predicted probability, connected by the solid line; the dashed line indicates perfect calibration.

The rolling-origin evaluation, which trains on all data through year Y and predicts year $Y + 1$, tells a more nuanced story (Table 6). Prospective accuracy is highly unstable: it reaches 0.79 in 2022 and 0.71 in 2023 but falls to 0.45 in 2021 and 0.43 in 2024—at or below chance—averaging 0.605 across the five test years. This instability, rather than the cross-validated mean, is the central empirical fact of the paper: the same model that ranks distress risk well in some years fails entirely in others. The variation does not track the number of events (2024 has the most events yet the weakest accuracy), pointing instead to a change in the underlying signal-to-distress relationship across years—the regime-dependence we examine next. A value below 0.50 in an out-of-sample evaluation deserves a precise reading: it does not merely indicate noise, but that the predictive relationships learned in the training window are locally inverted in the test window—coins scored as safer were, on balance, marginally more likely to enter distress. In turbulent regimes, that is, the market is not simply blind to the standard signals; it can act in direct opposition to them, which is the sharpest form of the regime-dependence this paper documents. Figure 3 displays these year-by-year accuracies.

Table 6. Rolling-origin out-of-sample AUC by test year.

Test Year	Obs.	Events	AUC
2021	225	36	0.451
2022	166	71	0.791
2023	642	22	0.713
2024	825	207	0.433
2025	474	170	0.635
Mean			0.605




Figure 3. Out-of-sample predictive accuracy by test year (rolling-origin). Bars below the dashed chance line (0.5), shown in red, indicate years in which the model performs no better than random. The instability—accuracy ranging from 0.43 to 0.79—is the central empirical fact motivating the regime analysis; the mechanism behind it is examined through the interaction model in Section 5.4, not through year-level averages. The green dotted horizontal line marks the mean rolling-origin AUC (0.605).

Table 7 extends the rolling-origin evaluation to the full metric set. The additional measures reproduce, and sharpen, the instability visible in the ROC AUC. In the strong years the model beats every baseline on every metric: in 2022 the precision–recall AUC was 0.755 against a no-skill baseline of 0.394, and in 2023 it was 0.127 against 0.034—a 3.7-fold lift—with Brier scores below the uninformative forecast in both years. In the weak years the model adds no value on any metric: in 2021 the precision–recall AUC (0.178) barely exceeded the event rate (0.155) and the Brier score was worse than the uninformative forecast, whereas in 2024 the precision–recall AUC (0.223) fell below the event-rate baseline (0.248) and the Brier score again exceeded the null. The regime-dependence documented below is therefore not an artifact of the choice of evaluation metric.

Table 7. Rolling-origin evaluation: imbalance-robust metrics by test year.

Test Year	PR-AUC	PR Baseline	Brier	Brier (Null)	Precision	Recall	F1
2021	0.178	0.155	0.147	0.136	0.147	0.694	0.243
2022	0.755	0.394	0.297	0.315	0.730	0.667	0.697
2023	0.127	0.034	0.042	0.058	0.085	0.455	0.143
2024	0.223	0.248	0.222	0.205	0.201	0.244	0.220
2025	0.469	0.370	0.265	0.275	0.432	0.485	0.457

PR baseline is the test-year event rate (the PR-AUC of a no-skill classifier); Brier (null) is the Brier score of a constant forecast at the training event rate; precision, recall, and F1 are computed at a classification threshold equal to the training event rate.

5.4. Regime-Dependence

To test whether the instability reflects a systematic dependence on market conditions, Table 8 augments the model with interactions between the standardized coin-level predictors and contemporaneous market-wide volatility (Equation (10)). The interactions are jointly highly significant: a likelihood-ratio test against the baseline yielded a chi-squared statistic of 23.7 on four degrees of freedom ($p \approx 9.1 \times 10^{-5}$), and the Akaike Informa-

tion Criterion fell from 2182 to 2167. Regime-conditioning therefore improves the model materially, not marginally.

Table 8. Predictor $\times$ market-volatility interactions. Inference by cluster bootstrap ($B = 1000$); see Section 6.4.

Interaction Term	Coefficient	Boot. p -Value
Realized vol. $\times$ market vol.	−0.122	0.078
Log volume $\times$ market vol.	−0.211	0.002
Quarterly return $\times$ market vol.	+0.177	<0.001

The interactions are individually interpretable. The volatility interaction was negative (−0.12) but only marginally significant under the conservative cluster-bootstrap inference ($p = 0.078$; see Section 6.4), a fragility we document carefully in robustness. The volume interaction was negative, the strongest of the three, and robust (−0.21, bootstrap $p = 0.002$), indicating that the liquidity signal loses force in turbulent regimes. The momentum interaction was positive and robust (+0.18, bootstrap $p < 0.001$), implying that the relationship between recent returns and distress shifts with the regime. Taken together, the coefficients show that the cross-sectional mapping from coin-level signals to distress is not fixed but conditional on the prevailing market state—with the volume and momentum channels well-identified and the volatility channel less so—which accounts for the year-to-year swings in out-of-sample accuracy documented above.

5.5. Mechanism

A natural explanation for the regime-dependence is that coins co-move more strongly when markets are turbulent, compressing the cross-sectional dispersion that coin-level signals exploit. We test this directly by relating the average pairwise return correlation across actively-traded coins (Equation (12)) to the contemporaneous volatility regime. The data do not support the channel: across 34 well-populated quarters the association between average co-movement and market-wide volatility is statistically indistinguishable from zero (Pearson $r = 0.08$, $p = 0.64$; Spearman $\rho = -0.02$, $p = 0.93$). Average cross-coin correlation in fact declines over the sample as the market broadens, a trend orthogonal to the volatility regime. The co-movement hypothesis is therefore rejected, and the mechanism underlying the regime-dependence remains open. We report this negative result transparently: ruling out the most obvious explanation, which sharpens the empirical regularity rather than weakening it.

6. Robustness

We subject the two headline results—modest out-of-sample predictability and significant regime-dependence—to a battery of robustness checks spanning the definition of distress, the choice of regime proxy, the data frequency, and the method of inference. The consistent finding is that both headline results are robust across every dimension, whereas the decomposition of the regime effect into individual signal channels is not. We report each in turn.

6.1. Sensitivity to the Distress Threshold

Our primary results use a drawdown threshold of $\theta = 0.70$. We re-estimated the entire pipeline at $\theta = 0.60$ and 0.80 (Table 9). Lowering the threshold mechanically raises the event rate (from 12.6% at 0.80 to 31.0% at 0.60) and enlarges the at-risk set, so the specifications overlap rather than nest. Out-of-sample accuracy was essentially invariant (cross-validated AUC 0.701, 0.678, 0.669), and the aggregate regime-dependence was significant at every

threshold (LR p from 2.5×10^{-8} to 6.2×10^{-9}). The composition, however, is threshold-sensitive: the momentum interaction is significant throughout, whereas the volatility and volume interactions fade at 0.80, where the distressed set is dominated by coins in outright collapse whose behavior may be less regime-contingent.

Table 9. Robustness across drawdown thresholds.

Threshold	Events	CV AUC	LR p -Value	Sig. Channels
60%	587	0.701	2.5×10^{-8}	vol, return
70% (primary)	529	0.678	9.1×10^{-5}	all three
80%	411	0.669	6.2×10^{-9}	return only

6.2. Alternative Regime Proxies

To confirm that the regime result does not depend on our particular measure of market conditions, we replaced the cross-sectional mean volatility with two Bitcoin-based proxies: Bitcoin's own drawdown and Bitcoin's realized volatility (Table 10). The aggregate regime-dependence survives both (LR $p = 0.029$ and 0.003 respectively), establishing that the finding is not an artifact of the chosen proxy. Consistent with Section 6.1, however, the channel that carries the effect differs across proxies—volume under the drawdown proxy, momentum under the volatility proxy—reinforcing our view that the aggregate effect is robust, whereas the per-channel decomposition is not.

Table 10. Regime-dependence under alternative market-regime proxies.

Regime Proxy	LR p -Value	Robust Channels
Cross-sectional mean vol. (primary)	9.1×10^{-5}	all three
Bitcoin drawdown	0.029	volume
Bitcoin volatility	0.003	return

6.3. Trailing-Peak Window and Data Frequency

The distress definition uses a 365-day trailing peak, and the panel is quarterly. We vary both (Table 11): trailing peaks of 180 and 540 days, and a monthly rather than quarterly frequency. Out-of-sample AUC remains in a narrow band (0.665–0.719) and the regime LR test remains significant in every case. The monthly specification is a particularly demanding check, tripling the sample to roughly 8000 observations and 860 events; the regime-dependence holds (LR $p = 0.013$), confirming the result is not an artifact of quarterly aggregation.

Table 11. Robustness across trailing-peak windows and data frequency.

Specification	Obs.	CV AUC	Regime LR p
Peak = 180 days, quarterly	3410	0.719	0.004
Peak = 365 days, quarterly (primary)	2579	0.680	9.1×10^{-5}
Peak = 540 days, quarterly	2196	0.665	1.8×10^{-5}
Peak = 365 days, monthly	7993	0.670	0.013

6.4. Clustered Inference via the Bootstrap

Analytic coin-clustered standard errors are numerically unstable in our setting because the combination of year fixed effects and clustering produces near-singular cluster contributions for coins observed in a single year. Therefore, we obtained clustered inference through a cluster bootstrap: we resampled coins with replacement (each drawn coin contributing all its coin-quarters), refit the interaction model, and constructed standard

errors and percentile confidence intervals from $B = 1000$ replications (all of which converged). Table 12 reports the interaction terms. The volume and momentum interactions remain firmly significant under this conservative inference (bootstrap $p = 0.002$ and <0.001 , with confidence intervals excluding zero). The volatility interaction weakens to marginal significance (bootstrap $p = 0.078$, confidence interval $[-0.242, 0.014]$), consistent with its fragility across the threshold and regime-proxy checks. Proper clustered inference thus confirms the aggregate regime-dependence while isolating the volatility channel as the least robust element.

Table 12. Regime interactions under cluster-bootstrap inference ($B = 1000$, resampling coins).

Term	Estimate	Boot. p	95% CI
Realized vol. $\times$ market vol.	−0.122	0.078	$[-0.242, 0.014]$
Log volume $\times$ market vol.	−0.211	0.002	$[-0.400, -0.071]$
Quarterly return $\times$ market vol.	+0.177	<0.001	$[0.082, 0.292]$

6.5. Robustness to the Predictive Method

The analyses above hold the predictive method fixed at the panel logit, which raises the question of whether the regime-dependence is a property of the market or of that functional form. We therefore repeated the out-of-sample evaluation with three flexible machine-learning classifiers—a random forest, gradient boosting, and XGBoost—trained on the identical predictor set and an at-risk sample and evaluated under the identical coin-stratified cross-validation and rolling-origin schemes (Table 13).

Table 13. Out-of-sample accuracy (ROC AUC) across predictive methods.

Model	CV AUC	2021	2022	2023	2024	2025	Rolling Mean
Panel logit (primary)	0.678	0.451	0.791	0.713	0.433	0.635	0.605
Random forest	0.740	0.502	0.715	0.739	0.456	0.582	0.599
Gradient boosting	0.739	0.525	0.714	0.717	0.475	0.610	0.608
XGBoost	0.740	0.473	0.742	0.672	0.464	0.614	0.593

CV is coin-stratified five-fold cross-validation; year columns are rolling-origin (train through year $Y - 1$, test on year Y). Random forest: 500 trees, minimum leaf size 10; gradient boosting: 300 trees, depth 3, learning rate 0.05, subsampling 0.8; XGBoost: 400 trees, depth 3, learning rate 0.05, row and column subsampling 0.8.

The flexible learners achieved somewhat higher cross-sectional discrimination (cross-validated AUC near 0.74, against 0.68 for the logit), consistent with the tree-based results of Będowska-Sójka et al. (2026). Their prospective accuracy, however, was exactly as unstable: all three failed in 2021 (AUC 0.47–0.52) and in 2024 (0.46–0.48) and succeeded in 2022–2023 (0.67–0.74), with rolling means of 0.59–0.61, indistinguishable from the logit. The regime-dependence is therefore a property of the market rather than of the panel-logit specification, and greater model flexibility does not restore reliability in turbulent regimes.

6.6. Researcher Degrees of Freedom and Multiple Testing

The analysis explored several specifications—three distress thresholds, alternative liquidity proxies, three regime proxies, two data frequencies, and alternative trailing-peak windows. To guard against selective reporting, we report results across all of these rather than a single preferred specification and frame the variation as robustness rather than selection. The two central findings are not the product of specification search: out-of-sample predictive accuracy is stable across every specification (cross-validated AUC 0.665–0.719) and the aggregate regime-dependence is significant at $p < 10^{-4}$ in nearly every case—a margin that survives any standard multiple-comparison adjustment, including a conservative Bonferroni correction across the full set of specifications. We are, by contrast, deliberately

cautious about the decomposition of the regime effect into individual channels, which is not stable: we treat only the volume and momentum channels, which survive conservative cluster-bootstrap inference, as well-identified, and explicitly flag the volatility channel as fragile.

6.7. Cross-Exchange Replication

The most consequential robustness question for a single-venue study is whether the finding is an artifact of the exchange. We therefore replicated the analysis on Binance, the largest cryptocurrency exchange by volume, whose spot markets are quoted in the USDT stablecoin rather than fiat USD. Using daily data for 79 USDT-quoted coins (2017–2026), we rebuilt the distress panel and re-estimated the regime model with the identical procedure (Table 14). The replication is demanding in two respects: it changes both the venue and the quote asset, and—because the Binance sample is drawn from currently-trading pairs—it is survivorship-biased toward survivors, the opposite of our survivorship-inclusive Kraken design.

Table 14. Cross-exchange replication: Kraken (primary) versus Binance.

Quantity	Kraken (Primary)	Binance (Replication)
Exchange/quote asset	Kraken/USD	Binance/USDT
Coins	609	79
At-risk coin-quarters	~2580	1370
Out-of-sample CV AUC	0.678	0.652
Regime-dependence LR p	9.1×10^{-5}	3.4×10^{-6}
Significant channels	volume, momentum	volatility, momentum

The finding holds. Out-of-sample predictive accuracy is comparable (cross-validated AUC 0.652 on Binance versus 0.678 on Kraken), and the aggregate regime-dependence is strongly significant (likelihood-ratio $p = 3.4 \times 10^{-6}$, against 9.1×10^{-5} on Kraken). As on Kraken, the aggregate effect is robust, whereas the contributing channels differ across the two exchanges—volatility and momentum are the significant interactions on Binance, volume and momentum on Kraken—independently reinforcing our central distinction between a robust aggregate effect and a specification-sensitive decomposition. That the regime-dependence survives a change of exchange, quote asset, and survivorship regime is strong evidence that it is a feature of cryptocurrency markets rather than of any single venue.

6.8. Summary

Across five independent dimensions—distress threshold, regime proxy, peak window and data frequency, method of inference, and a second exchange—the two headline results hold: distress is modestly predictable out-of-sample and the predictability is regime-dependent. The single consistent qualification is that the volatility channel, and more generally the decomposition of the regime effect into individual channels, is sensitive to specification; the aggregate effect is not. We regard the cross-exchange replication as substantial evidence that the regime-dependence is a genuine property of cryptocurrency distress dynamics rather than a feature of the Kraken sample.

7. Discussion

7.1. Interpretation

Our findings refine, rather than overturn, the emerging consensus that cryptocurrency distress is predictable from market data. We confirm that microstructure indicators—volatility, liquidity, momentum, and age—carry information about subsequent distress,

reproducing the broad thrust of recent failure-prediction studies. The novel and, we think, more consequential finding is that this predictability is conditional: the very signals that flag vulnerable coins in calm markets lose much of their discriminating power precisely when market-wide volatility is elevated—that is, when distress is most widespread and an early warning would be most valuable. The model is most reliable when it is least needed and least reliable when it is most needed. This asymmetry is the core economic content of the paper.

7.2. *Connection to the Volatility-Regime Literature*

The result links two literatures that have largely developed separately: the prediction of crypto asset failure and the study of time-varying volatility and connectedness in crypto markets. The latter establishes that cross-asset volatility structure intensifies and reorganizes during turbulent periods. It is therefore coherent that the mapping from coin-level signals to distress should also shift with the volatility regime. Our contribution is to demonstrate this conditioning directly within a predictive framework and to quantify its cost in out-of-sample accuracy, rather than to infer it indirectly from the volatility process alone.

7.3. *A Tested Mechanism, Not Supported*

A natural explanation for the regime-dependence is that coins co-move more strongly when market-wide volatility is high, compressing the cross-sectional dispersion that coin-level signals exploit and thereby degrading predictability. Section 5.5 tests this channel directly and finds no support for it: average cross-coin correlation in fact declines over the sample as the market broadens, a trend orthogonal to the volatility regime. The co-movement mechanism is therefore rejected, and the source of the regime-dependence remains an open question. We report this negative result transparently; ruling out the most obvious mechanism sharpens, rather than weakens, the regularity we document.

The rejection of the co-movement channel invites consideration of alternatives, which we outline as hypotheses for future work. First, turbulent regimes may compress the cross-sectional dispersion of liquidity rather than of returns, blunting the volume signal specifically—a reading consistent with the volume interaction being the strongest and most robust channel in Table 8. Second, sentiment- and attention-driven flows may dominate coin-specific fundamentals in high-volatility episodes, temporarily decoupling microstructure signals from distress risk; the documented informativeness of online attention for coin survival (Fantazzini, 2022) makes this channel both plausible and testable with attention data. Third, flight-to-quality toward the largest coins during stress may reorganize the trading of small coins in ways that average pairwise correlation does not capture, suggesting concentration- or dominance-based regime measures. Fourth, regime shifts that coincide with macro-financial or regulatory shocks may change which coins are vulnerable rather than merely how vulnerable they are, breaking any fixed cross-sectional mapping. Finally, exchange-level listing and delisting policy may itself be regime-sensitive, although the cross-exchange replication of Section 6.7 suggests the phenomenon is not venue-specific. Discriminating among these candidates requires data beyond the OHLCVT records used here—sentiment and attention indices, dominance measures, and uncertainty proxies—and is a natural next step for this agenda. We note also a methodological refinement of the test itself: the co-movement channel was assessed here through unconditional quarter-level average correlations, and dynamic conditional correlation models (e.g., DCC-GARCH) would permit a sharper, time-varying test of the same hypothesis—capturing the possibility that correlations spike non-linearly within turbulent quarters in ways that quarterly average obscures.

7.4. Implications

For practitioners, the central lesson is that early-warning tools built on coin-level microstructure features should not be assumed to perform uniformly. Their reliability is itself state-dependent; risk systems that rely on such signals should condition on—or at minimum stress-test against—the prevailing volatility regime and should attach lower confidence to distress forecasts generated in turbulent markets. For exchanges, the same conditionality bears on the design of listing-monitoring and delisting-review processes, which may need regime-aware thresholds rather than fixed ones. For investors and those concerned with investor protection, the result tempers optimism that distress can be reliably anticipated from public market data alone: it can, but only sometimes, and not in the conditions where protection matters most.

These implications can be made operational. For portfolio and risk managers, distress scores from microstructure models should carry regime-indexed confidence: score-based position limits can be widened, and automated de-risking triggers suspended or supplemented with discretionary review, whenever market-wide volatility exceeds a calibrated bound—in our rolling evaluation, forecasts issued in such periods were no better than chance (AUC 0.43–0.45) and, by imbalance-robust measures, added no value at all (Table 7). For exchanges, surveillance thresholds calibrated in calm markets will systematically over- or under-flag coins in turbulent ones; a two-tier review that increases reliance on non-market information—development activity, disclosure quality, on-chain usage—when the volatility regime is elevated follows directly from our estimates. For regulators and investor-protection bodies, the finding cautions against warning systems that rely on public market data alone: such systems cannot be presumed to function during systemic crypto stress, precisely the episodes of greatest concern, and should be complemented with structural information sources. In each case, the operative principle is the same: both the forecast and the confidence attached to it must be conditioned on the regime.

Practitioners will reasonably ask where the regime boundary lies. Our estimates support an approximate rather than a sharp answer. In our sample the regime variable—the cross-sectional mean of annualized realized volatility across actively traded coins—has a long-run mean near 1.1 with a standard deviation of roughly 0.3. Quarters more than about one standard deviation above that mean (cross-sectional mean volatility above roughly 1.4) are the neighborhood in which 2021, the clearest failure year, sits, and the interaction estimates imply that the return and volume channels are largely neutralized at such levels. As a working rule of thumb, we would treat model scores with suspicion whenever trailing-quarter market volatility exceeds its historical 75th–80th percentile—about half a standard deviation above its long-run mean, the top quarter of quarters in our sample—and as unreliable beyond one standard deviation. Two qualifications matter. First, any fixed cut-off should be paired with a rolling out-of-sample check of the kind in Section 5.3, since 2024 demonstrates that prospective accuracy can deteriorate even when average volatility is only moderately elevated: the volatility regime is a warning indicator, not a sufficient statistic. Second, the percentile itself should be recalibrated as data accrue, since the level of market-wide volatility has drifted downward as the market has matured.

7.5. Limitations

Several limitations bound our conclusions and define the agenda for future work. Our primary distress measure was observed on a single exchange; although we replicated the regime-dependence on Binance (Section 6.7), both venues captured venue-specific rather than global distress, and a coin inactive on one may trade on another. The analysis is quarterly, which suits a definition of sustained distress but coarsens the timing of onset, though our monthly robustness specification mitigates this concern. The mechanism behind

the regime-dependence is, as noted, unresolved. Although the aggregate results are robust across thresholds, regime proxies, frequencies, inference methods, and exchanges, the per-channel decomposition is not, with the volatility channel in particular sensitive to specification. We regard these as boundaries on the present claims rather than threats to them. Two further caveats bear stating. Our estimates are predictive associations, not causal effects: volatility and volume are jointly determined with incipient deterioration, and omitted fundamentals plausibly drive both predictors and outcome (Section 4.4). Although the pooled hazard specification is deliberate—our object is prediction for previously unseen coins—estimator choice does not drive the headline result, which is unchanged under GEE and random-effects re-estimation.

8. Conclusions

This paper asks whether market-microstructure indicators predict the onset of severe cryptocurrency distress, and whether that predictability is stable across market conditions. Using the full available history of 609 USD-quoted coins on the Kraken exchange and a survivorship-inclusive coin-quarter panel, we found that distress onset is modestly predictable out-of-sample—a coin-stratified cross-validated AUC of about 0.68—from volatility, liquidity, momentum, and age, confirming prior work.

Our central contribution is to show that this predictability is regime-dependent. Interactions between coin-level signals and market-wide volatility are jointly and individually significant, a rolling-origin evaluation reveals prospective accuracy ranging from no better than chance to strong across years, and the aggregate regime-dependence is robust to the definition of distress. We further tested the most natural explanation—rising cross-asset co-movement in turbulent markets—and rejected it, leaving the mechanism as an open question.

The message for researchers and practitioners alike is that early-warning signals in cryptocurrency markets are conditionally reliable, and that treating them as time-invariant overstates their practical value, particularly in the turbulent conditions where reliable warning is most needed. Future work should identify the mechanism behind the regime-dependence, extend the analysis across multiple exchanges and frequencies, and investigate whether explicitly regime-aware models can recover stable predictive performance. Establishing that predictability is conditional is, we hope, a useful step toward building tools that account for that conditionality rather than ignoring it.

Author Contributions: Conceptualization, A.A. and H.A.; methodology, H.A.; software, H.A.; resources, A.A.; validation, H.A.; formal analysis, H.A.; data curation, H.A.; writing—original draft preparation, H.A.; writing—review and editing, A.A. and H.A.; supervision, A.A.; funding acquisition, A.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Deanship of Scientific Research at Imam Mohammad Ibn Saud Islamic University (IMSIU) through the Deputyship of Scientific Research and Publication, grant number IMSIU-DDRSP2604.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data used in this study are derived from publicly available exchange records (Kraken historical OHLCVT data and Binance public market data). Processed datasets and analysis code are available from the authors on request.

Acknowledgments: The authors gratefully acknowledge the support of the Deanship of Scientific Research at Imam Mohammad Ibn Saud Islamic University (IMSIU) through the Deputyship of Scientific Research and Publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1), 31–56. [\[CrossRef\]](#)
- Ammann, M., Burdorf, T., Liebi, L., & Stöckl, S. (2022). *Survivorship and delisting bias in cryptocurrency markets* (Working paper). University of St. Gallen and University of Liechtenstein. [\[CrossRef\]](#)
- Antonakakis, N., Chatziantoniou, I., & Gabauer, D. (2019). Cryptocurrency market contagion: Market uncertainty, market complexity, and dynamic portfolios. *Journal of International Financial Markets, Institutions and Money*, 61, 37–51. [\[CrossRef\]](#)
- Baruník, J., & Křehlík, T. (2018). Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics*, 16(2), 271–296. [\[CrossRef\]](#)
- Będowska-Sójka, B., Wójcik, P., & Pele, D. T. (2026). Early warning systems for cryptocurrency markets: Predicting ‘zombie’ assets using machine learning. *The North American Journal of Economics and Finance*, 81, 102543. [\[CrossRef\]](#)
- Bianchi, D., Guidolin, M., & Pedio, M. (2023). The dynamics of returns predictability in cryptocurrency markets. *The European Journal of Finance*, 29(6), 583–611. [\[CrossRef\]](#)
- Brauneis, A., Mestel, R., Riordan, R., & Theissen, E. (2021). How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance*, 124, 106041. [\[CrossRef\]](#)
- Buse, R., Görgen, K., & Schienle, M. (2025). Predicting value at risk for cryptocurrencies with generalized random forests. *International Journal of Forecasting*, 41(3), 1199–1222. [\[CrossRef\]](#)
- Campbell, J. Y., Hilscher, J., & Szilagyi, J. (2008). In search of distress risk. *The Journal of Finance*, 63(6), 2899–2939. [\[CrossRef\]](#)
- Chava, S., & Jarrow, R. A. (2004). Bankruptcy prediction with industry effects. *Review of Finance*, 8(4), 537–569. [\[CrossRef\]](#)
- Dangl, T., & Halling, M. (2012). Predictive regressions with time-varying coefficients. *Journal of Financial Economics*, 106(1), 157–181. [\[CrossRef\]](#)
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66. [\[CrossRef\]](#)
- Fantazzini, D. (2022). Crypto coins and credit risk: Modelling and forecasting their probability of death. *Journal of Risk and Financial Management*, 15(7), 304. [\[CrossRef\]](#)
- Fantazzini, D., & Calabrese, R. (2021). Crypto exchanges and credit risk: Modeling and forecasting the probability of closure. *Journal of Risk and Financial Management*, 14(11), 516. [\[CrossRef\]](#)
- Fantazzini, D., & Zimin, S. (2020). A multivariate approach for the simultaneous modelling of market risk and credit risk for cryptocurrencies. *Journal of Industrial and Business Economics*, 47(1), 19–69. [\[CrossRef\]](#)
- Feder, A., Gandal, N., Hamrick, J. T., & Moore, T. (2018, June 18–19). *The rise and fall of cryptocurrencies*. Workshop on the Economics of Information Security (WEIS), Innsbruck, Austria.
- Figà-Talamanca, G., Focardi, S., & Patacca, M. (2021). Regime switches and commonalities of the cryptocurrencies asset class. *The North American Journal of Economics and Finance*, 57, 101425. [\[CrossRef\]](#)
- Gandal, N., Hamrick, J. T., Moore, T., & Vasek, M. (2021). The rise and fall of cryptocurrency coins and tokens. *Decisions in Economics and Finance*, 44(2), 981–1014. [\[CrossRef\]](#)
- Henkel, S. J., Martin, J. S., & Nardari, F. (2011). Time-varying short-horizon predictability. *Journal of Financial Economics*, 99(3), 560–580. [\[CrossRef\]](#)
- Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters*, 158, 3–6. [\[CrossRef\]](#)
- Konuk, D. E., & Güvenir, H. A. (2025). Predicting the risk of death for cryptocurrencies using deep learning. *Journal of Risk and Financial Management*, 18(12), 716. [\[CrossRef\]](#)
- Kraken. (n.d.). *Downloadable historical OHLCVT data*. Kraken. Available online: <https://support.kraken.com/articles/360047124832-downloadable-historical-ohlcvt-open-high-low-close-volume-trades-data> (accessed on 1 May 2026).
- Lansky, J. (2020). Cryptocurrency survival analysis. *The Journal of Alternative Investments*, 22(3), 55–64. [\[CrossRef\]](#)
- Lee, G., & Ryu, D. (2026). State-dependent relationship between cryptocurrency returns and credit spreads. *European Financial Management*, 32, 261–287. [\[CrossRef\]](#)
- Maciel, L. (2021). Cryptocurrencies value-at-risk and expected shortfall: Do regime-switching volatility models improve forecasting? *International Journal of Finance & Economics*, 26(3), 4840–4855. [\[CrossRef\]](#)
- Pettenuzzo, D., Timmermann, A., & Valkanov, R. (2014). Forecasting stock returns under economic constraints. *Journal of Financial Economics*, 114(3), 517–553. [\[CrossRef\]](#)
- Rapach, D. E., Strauss, J. K., & Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies*, 23(2), 821–862. [\[CrossRef\]](#)

- Shumway, T. (2001). Forecasting bankruptcy more accurately: A simple hazard model. *The Journal of Business*, 74(1), 101–124. [[CrossRef](#)] [[PubMed](#)]
- Yang, S., & Kwon, H. (2026). Developing predictive and explainable models for cryptocurrency delistings: A case study of Binance exchange. *Asia-Pacific Journal of Financial Studies*. Advance online publication. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.